Subject: Re: Bizarre slowness from sort()
Posted by William Clodius on Tue, 23 Apr 2002 21:11:32 GMT
View Forum Message <> Reply to Message

Jonathan Joseph wrote:

> I had the following enlightening discussion about qsort with a
> programmer from RSI who responded to my post, but I'm still not
> entirely convinced. Can anyone comment on this?
>
> -Jonathan
>
> (posted with permission)<snip>

They should not be using quicksort in any form for integer or floating point data. The quoted average 0(n ln(n)) performance is not as good as the 0(n) that is achievable for these special data types. (Further this "average" has assumptions about the randomness of data that is often not obeyed in practice.) For other datatypes, e.g., strings, they should also not be using quicksort, as other sorting algorithms have the same 0(n ln(n)) asymptotic behavior and their relative scale factors depend on the relative costs of memory fetching, storing, and comparisons. Typical arguments favoring quicksort are based on values for these relative costs that were typical for integer and floating point data on machines of the late 70's/early 80's. See Richard O'Keefe's discussion of his applicative merge sort in a variety of newsgroups in the 90's.